CHAPTER

3

INSTALLING LINUX

After reading this chapter and completing the exercises, you will be able to:

- Discuss issues related to installing Linux
- Install a popular distribution of Linux
- Start using a new Linux system
- Troubleshoot problems with a Linux installation

In the previous chapter you learned how to gather information about hardware components and networking protocols as preparation for installing Linux. You also learned about preparing a computer running Microsoft Windows to be used as a dual-boot system with Linux.

In this chapter you will learn how to prepare for and install a new Linux operating system. Among other things, you will learn how to prepare Linux partitions, configure the system, and answer questions during the installation process. You will also learn what to do when the installation process doesn't work correctly. Finally, you will learn about starting Linux for the first time, after completing the installation.

UNDERSTANDING INSTALLATION ISSUES

The following sections review a few general issues related to installing Linux and explain how to answer the specific questions that arise during the installation process. Although the later parts of this chapter discuss installing the Red Hat Linux distribution included with this book, you may choose to install and use several different Linux products as you learn about Linux. Thus, this discussion concerns issues that are common to all versions of Linux.

Linux Distributions

People seem to like having favorites: one person insists that Ford trucks are better; another buys only Chevrolet. One person always drinks Coke; another chooses Pepsi every time. The differences between products may be slight, but the loyalty they inspire is not. Linux users often have similar feelings about the version or distribution they have chosen. One person insists that Debian Linux is the only reasonable choice; another uses Red Hat exclusively.

Although having a favorite is fine, remember that Linux distributions actually are very similar to each other. Each one takes the Linux kernel from the same location on the Internet, and each uses the same set of supporting utilities. Table 1-1 (in Chapter 1) names several of the better-known distributions. The list below highlights the major differences among Linux distributions. As this list indicates, you can choose a distribution with characteristics that fit your preferences.

- The installation program for each distribution is different. Linux vendors put a lot of time and money into designing their installation programs, in order to meet the needs of particular audiences. Thus, some installation programs focus on ease of use for those who are new to Linux; others focus on flexibility for users who are very familiar with Linux features.
- The arrangement of the files that make up a standard Linux system is fairly standard, but some variations are possible. Different vendors place certain configuration files in different locations. (These differences can be traced to the preferences of the software developers who assemble a particular distribution—each group of software developers has its own idea of the best way to arrange the files.) For example, Caldera OpenLinux places configuration files for the Samba network service in the directory `/etc/samba.d`, whereas Red Hat Linux places these files in the directory `/etc`. The variety of locations for configuration and initialization files makes moving from one Linux distribution to another a challenge. But overall, the differences are small compared to the total number of files installed for a typical Linux system.
- The default configuration values and the default services started on a new Linux system vary by distribution. One vendor might choose to start services or set up a user's environment slightly differently from another vendor. You can change these options after installing Linux.
- Some configuration tools used in a distribution are specifically developed by the vendor of that distribution. You will learn more about these tools in later chapters (beginning in Chapter 7). The text configuration files that control Linux are the same for all distributions, but the graphical utilities available for configuring these files are sometimes specific to one product. For example, only Red Hat Linux includes the LinuxConf utility; only Suse Linux includes the YAST utility.
- In general, each Linux vendor targets a certain type of customer. The intended customer might be very technical or very new to Linux; the customer might be business oriented or a home user; the customer might be in the United States or in another country. Linux vendors have wisely decided not to attempt to be all

things to all people. Hence the distribution that you choose should match the type of customer you are (or that you intend to work for).

- Certain Linux vendors focus on creating products with the most recent versions of all software from the Internet, but they must sacrifice the time-consuming testing and documentation that other vendors choose to invest in. Having the very latest Linux features as soon as possible is a big concern for some Linux users; others prefer software that has been more thoroughly prepared for stable long-term use. For example, Debian Linux is generally considered to be dedicated to free software ideals; Red Hat Linux is usually the most current with recent Linux releases; Caldera OpenLinux is the most thoroughly tested and stable distribution; and SuSE Linux is the most suitable for international use.

Red Hat Linux

A few words about the Red Hat Linux distribution are in order here. Red Hat Linux is the most popular Linux distribution in the world (speaking of the number of people using the product). It has been around since about 1993. Red Hat was created for technical people who are actively following the Linux development community on the Internet, although more recently, Red Hat has been focusing on business users and ease of use. Some people prefer other distributions over Red Hat, but because of its popularity in the marketplace, Red Hat was selected for inclusion in this book as an installation example. Notes about installing other distributions are also included in this chapter.

The Red Hat Linux CD included with this book is the basis for the installation instructions provided in this chapter. You can, however, obtain a copy of several versions of Linux in any of the following ways:

- Use the CD provided with another book on Linux.
- Receive a free copy of Linux on CD at a trade show or other event sponsored by a Linux vendor.
- Purchase a Linux CD for less than $5 from a company such as Linux Mall (at *www.linuxmall.com*).
- Purchase a retail copy of Linux at a software store or bookstore. (Retail versions often include multiple CDs and a printed user manual.)
- Download Linux from the FTP site of a Linux vendor, such as *ftp.redhat.com*. (This option is free but takes a long time to download unless you have a very fast Internet connection.)

The Installation Process

Installing an operating system on a computer is different from installing an application like a word processor. When you install an application, the existing operating system takes control and provides a foundation for the installation process. When you install a new operating system, only the hardware is available—no other software can assist the installation process.

The new operating system must somehow initialize itself sufficiently to install itself on the computer.

You use a Linux installation utility to install Linux onto the hard disk of a computer. The general procedure is as follows:

1. Start the installation program from a CD, floppy disk, or other source.
2. Run a very small copy of Linux within the computer's RAM.
3. Determine (based on user input or by automatically probing the system) where the installation source data is located. (This is described in the next section.)
4. Determine (based on user input or by automatically probing the system) where the Linux operating system should be installed. A **target hard disk partition**, or target partition, is the location on the system's hard disk where Linux will be installed.
5. Collect user input to determine which software packages to install on the target partition.
6. Collect user input to determine configuration settings such as network addresses, preferred services to execute, keyboard settings, and so forth.
7. Copy the Linux software packages from the installation source to the target partition.
8. Finish setting up the software packages on the target partition so they are ready to use.
9. Configure the LILO boot manager and install it on the boot sector (or MBR) of a hard disk so that the computer's BIOS can launch Linux.
10. Restart the computer to launch the newly installed Linux operating system.

Some of the steps in the previous list are handled automatically by the installation program you use to install Linux; most steps require that you answer one or more questions about how the Linux system will be set up. More detailed information about these steps is presented later in this chapter. This information will help you understand how to answer the questions presented during a Linux installation.

Installation Source Options

The **installation source** is the set of files from which Linux is installed. In this book, it is assumed that your files are stored on a Linux CD and then copied to the hard disk in an orderly manner during the installation process. If your installation files are stored on another medium, you can specify the exact location as you install Linux. The following list summarizes the most common installation sources:

- Local CD-ROM: use the files from the CD-ROM drive attached to the system on which you are installing Linux.

- Local hard disk: use a copy of the CD-ROM files that has been stored on a hard disk within the system on which you are installing Linux. (This must be a hard disk partition or hard disk that is separate from the location where Linux will be installed.)
- Floppy disks: read a series of floppy disks in order to retrieve all the Linux installation source files during the installation process. This method is rarely used for standard Linux distributions because it would require over 100 disks for most installations. Also, CD-ROM drives are very commonly available now.
- Network installation: read the Linux installation source files from a server connected to the same network as the computer on which you are installing Linux. You can use the network to access installation data in a variety of formats, including FTP (a standard Internet format), SMB (used by Microsoft Windows servers), and NFS (the Network File System protocol, commonly available on all Linux and UNIX systems). Using a network installation requires some additional configuration on the server that will provide the installation source, but this method allows you, as a system administrator, to install multiple Linux systems from a single networked copy of the installation source.

A few specialized versions of Linux are completely contained on a single floppy disk. For information on one of these specialized products, visit the homepage of the Linux Router Project at *www.linuxrouter.org*.

When installing Linux from a network, be careful not to violate copyright restrictions by installing a Linux distribution on systems for which you do not have a software license. Many Linux products have no restrictions on installing the product on more than one system; some commercial versions of Linux, or software included with commercial versions of Linux, may only permit you to install the software on one computer for each copy you have purchased.

To use a networked installation source, you must prepare another computer to provide the installation data. Because this involves setting up networking services that are beyond the scope of this book, you should contact the system administrator of the networked server you want to use. To use a Microsoft Windows server for the Linux installation source, ask the system administrator of the Windows system to prepare a share containing the CD-ROM drive (or a hard disk onto which the CD data has been copied). To use the NFS protocol on a Linux or UNIX server, ask the system administrator to export the CD or hard disk directory containing the Linux installation source. Some Linux vendors provide FTP servers for installations that download files from the Internet as the product is installed. (If the version of Linux you are installing supports FTP-based installation, you will see an option for selecting FTP during the installation program.)

Red Hat Linux supports installation from CD-ROM or local hard disk, plus network-based installations from NFS, SMB, or FTP servers.

Starting the Installation

When you turn on a computer, the BIOS normally sends control to the MBR of the first hard disk so that the operating system on the hard disk can launch. To install Linux, however, you must pass control directly to the Linux installation program without first starting the operating system on the hard disk. In the past, this was normally accomplished by using a **boot disk** (a floppy disk that has a portion of the Linux installation program on it). But newer computer systems have bootable CD-ROM drives. A **bootable CD-ROM drive** is one that can launch an operating system (or other program) directly from a CD without accessing the hard disk.

Most Linux CDs (including the one provided with this book) are bootable CDs. To see if your CD-ROM drive is bootable, try inserting the installation CD into the CD-ROM drive and then restarting the computer. If the installation program launches, your CD-ROM drive is bootable.

If the installation program does not start from the installation CD, it's possible that your CD-ROM drive is bootable, but that the BIOS is configured to start the operating system from the hard disk first. You can change the order in which devices are checked at boot time by reconfiguring the BIOS as described in Chapter 2.

Not all systems can boot directly from a CD. As a backup, virtually all systems can start the computer from a boot disk. (A boot disk is also called an **install disk** on some distributions of Linux.) After starting the system with a boot disk, the installation program locates the installation source on a CD or network server and proceeds with the installation. The CD containing the Linux installation source data also contains a copy of a boot disk in the form of a disk image. A **disk image** is a single file that contains an exact copy of a floppy disk. You can copy the disk image from the CD to a floppy disk by using the `rawrite` utility in DOS or Windows, or the `dd` utility on an existing Linux system.

Often you will need two disks to start the Linux installation because the necessary data will not fit on one disk. This is especially true if you are installing Linux on a laptop computer. Installing Linux on laptops is generally more challenging than installing it on a desktop system because laptops are more likely to contain unusual hardware components or components that are very new (and thus lack complete support in Linux). A great resource for information about installing Linux on various laptop models is the Linux on Laptops Web page, located at *www.cs.utexas.edu/users/kharker/linux-laptop/*.

Different distributions of Linux use different names for the disks used to start the Linux installation. One distribution calls the disks "Install" and "Modules," another calls them "boot" and "root." The Red Hat Linux CD included with this book uses the following three disk names. Each one is a file (with the file extension .img) located in the images subdirectory of the CD.:

- boot: used for standard installations from a CD.
- bootnet: used when the installation source is located on a network server (such as an NFS server).

- pcmcia: used to install Red Hat Linux on a laptop when the PCMCIA cards must be accessed to complete the installation. This occurs in two situations: when the device containing the installation source (such as a SCSI hard disk) is attached using a PCMCIA expansion card, or when installation is from a network server through a PCMCIA network adapter card. If neither of these situations applies, only the boot or bootnet disk is required.

To make a disk from a disk image on a CD in a Microsoft Windows computer, follow these steps. Note that these steps work for any Linux distribution, although the disk image names given in the steps are specific to the Red Hat Linux CD.

1. Insert the Red Hat Linux CD in a computer running Windows 95/98.
2. Use the Windows Explorer window to display the contents of the `dosutils` subdirectory on the CD. Note the drive letter associated with the CD-ROM drive (normally D: or E:).
3. Double-click the icon for the `rawrite` program. The program launches.
4. When you see the prompt for the name of a disk image file, enter the filename of the disk image you want to create. For example, if the CD–ROM drive is drive D: and you need to create the boot disk for Red Hat Linux, enter this:

   ```
   D:\images\boot.img
   ```

5. You are now prompted for the drive letter of the floppy disk drive. Enter the letter for this disk drive (normally either A: or B:).
6. You are prompted to press Enter to begin creating the disk. Press Enter to do so. The process of copying the image file to disk may take several minutes. When the process is completed, the floppy disk drive stops rotating. You can then close the `rawrite` program window in Windows and remove the disk you created.
7. Label the disk you created as **Red Hat Boot Disk** and put it in a safe place, so you can use it to install Linux. (Labeling the disk is especially important if you must use more than one disk to start Linux; otherwise, you might attempt to begin the installation with the wrong disk.)

THE INSTALLATION PROCESS

When you begin the Linux installation program, you must respond to a series of questions about how to install and configure the new Linux system. These questions may appear on a graphical screen, as shown in Figure 3-1, or on a text-based screen, as shown in Figure 3-2. Each version of Linux uses a different installation program, though they all ask very similar questions to complete the installation process. Some installation programs (such as the one for Slackware Linux) assume you know a lot about your hardware and about Linux. They ask many questions that a new Linux user may not be prepared to answer. Other products, such as Red Hat Linux and Caldera OpenLinux, are intended for new Linux users. They automate many functions and ask you fewer questions as you install Linux. Although some users might prefer to answer as few questions as possible, others appreciate the customization options such questions imply.

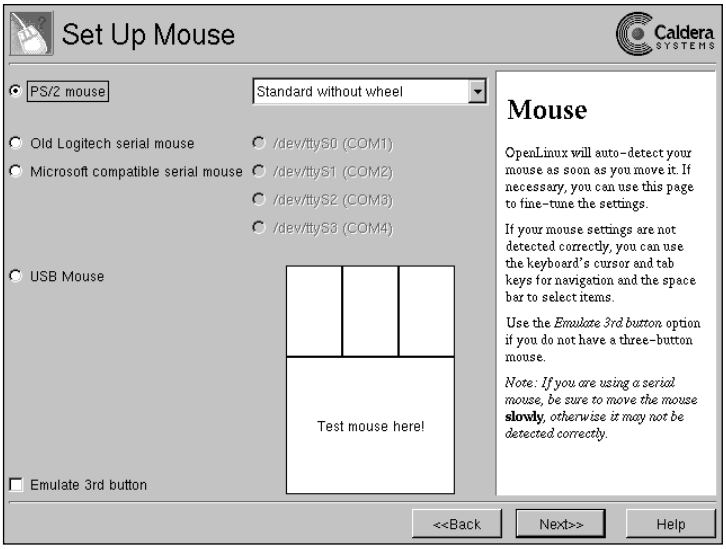

Figure 3-1 Typical graphical installation screen

Figure 3-2 Typical text-based installation screen

As you install a Red Hat Linux distribution, you have several opportunities to enter information about your system and select which components of Linux you want to install, but overall, Red Hat Linux tries to minimize the information needed from the person completing the installation. Red Hat Linux also provides both a graphical installation and a text-based installation. The graphical installation is used by default, but on many systems you may encounter problems with the video card. For these systems you can use the text-based installation program.

The sections that follow explain some of the information that you must provide during a typical Linux installation, with specific comments related to the Red Hat Linux product included with this book.

Projects 3-1 and 3-2, at the end of this chapter, provide step-by-step instructions for installing the version of Red Hat Linux provided with this book.

Answering Initial Questions

After you have started the installation program from a bootable CD-ROM or floppy drive, you need to answer questions about how to interact with the installation program. For example, you must choose a keyboard layout and a language for the installation. Depending on the version of Linux you are using, you might also be asked about your mouse, time zone, video card, or other details before selecting an installation source and target partition.

In many versions of Linux, you can back up to change your answer to previously asked questions. Because nothing is written to the hard disk of the computer until you have specified a target hard disk partition, you can also turn off the computer and restart the installation if you become concerned that you have not made correct selections. (Because nothing is written to hard disk in the early part of an installation, you can turn off the computer's power without the risk of losing data.)

Preparing Hard Disk Partitions

In Chapter 2 you learned how to prepare space on a Microsoft Windows computer so that you can install Linux on a hard disk where Windows is already installed. In this section you learn how to use that free space, or free hard disk space on any hard disk, to prepare for the installation of Linux.

The Linux installation program allows you to set up the partitions on your hard disk. Setting up Linux partitions involves the following steps, as you may recall from Chapter 2:

- Define a swap partition
- Define a Linux partition
- Mark the Linux partition as the active partition

Most Linux installation programs let you use a utility called `fdisk` (for *fixed disk*, meaning a hard disk). The **`fdisk`** utility is used to create partitions and configure how they are used. This utility is started from the installation program. (It can also be used after you have installed Linux to make changes in hard disk partitions.) Other utilities that have been developed by Linux vendors provide similar functionality, usually with an easier-to-use interface.

Hard Disk Specifications

Hard disks in Linux are identified by device names that resemble directory names. Later in this book you will learn more about how devices in Linux are accessed via the same directory structure as normal files. For now, you only need to know that certain names represent hard disk devices and partitions in Linux.

Hard disks that are attached to the computer using the IDE interface are identified as `/dev/hda` for the first hard disk, `/dev/hdb` for the second hard disk, and so forth, to `/dev/hdd` for the fourth hard disk. If you need to refer to a specific partition, you can add a partition number after the device name. For example, the first partition on the second IDE hard disk is represented as `/dev/hdb1`.

Hard disks that include the SCSI interface use a similar pattern of names, but with the letters *sd* instead of *hd*. For example, the first SCSI hard disk is `/dev/sda`. The second partition on the second SCSI hard disk is `/dev/sdb2`.

You will use these same device names to refer to a CD-ROM drive that is attached to the IDE or SCSI controller card on the system. For example, the CD-ROM drive is often attached as the third IDE device (the first device on the second IDE controller). Thus, the CD-ROM can be accessed by referring to `/dev/hdc`. (CD-ROMs do not have multiple partitions as hard disks do.)

Deciding on Mount Points

Unlike other operating systems, a Linux system does not use drive letters. Instead, different devices are accessed using subdirectories of a single directory structure. For example, in Windows, you might use the designations shown in Table 3-1.

Table 3-1 Typical Drive Letter Assignments in Windows

Drive	Description
A:	Floppy drive
C:	Main Windows hard disk partition
D:	Secondary Windows hard disk partition or second Windows hard disk
E:	CD-ROM drive
F:	Network server home directory

Instead of using drive letters, Linux defines a mount point for different file systems or devices that can contain data. A **mount point** is a subdirectory through which a set of data is accessed. Table 3-2 shows how the devices in Table 3-1 might be represented on a Linux system.

Table 3-2 Typical Subdirectory Mount Points Used to Access File Systems in Linux

Subdirectory (mount point)	Description
`/mnt/floppy`	Floppy drive
`/`	Main Linux hard disk partition
`/opt`	Secondary Linux hard disk partition or second Windows hard disk
`/mnt/cdrom`	CD-ROM drive
`/remote_home`	Network server home directory

The directory names used as mount points are arbitrary, but standard names are normally used for some devices, such as `/mnt/floppy` for the floppy disk drive. You will learn much more about setting up mount points in Linux in Chapter 8.

To complete the installation, you must be prepared to specify mount points for the new Linux system. The only mount point you *must* specify is the root partition, for the root directory, `/`. You must define which hard disk partition will hold this directory and its subdirectories. You can also define mount points for other hard disk partitions if you choose to. Because devices such as floppy disks and CD-ROM drives are not permanently available (you can insert and remove multiple disks as you work in Linux), these devices are not used during installation as mount points for installation files.

In Chapter 4 you will learn about the subdirectories that make up a standard Linux file system. One of the standard subdirectories is `/home`, where the home directory for each user account on the system is normally stored. For example, `/home/nwells` is a home directory for the user account `nwells`. During the installation you can place different subdirectories on different hard disk partitions by defining the subdirectory as a mount point.

When installing a very large Linux server supporting hundreds of users, it would be common to place the core operating system on one hard disk, all the user home directories on a second hard disk, and all the applications and data on a third hard disk. You would define this type of

configuration during the installation as you set up the mount points for various directories. In this example, you might define the mount points listed in Table 3-3.

Table 3-3 Example Mount Points for a Multiple Partition Installation

Mount point	Device
/	/dev/hda2
/home	/dev/hdb1
/opt	/dev/hdb2

Your first few Linux installations shouldn't require anything but the / mount point for the main Linux system. The entire operating system will then be installed on a single partition. The following list gives a few reasons why you might want to use multiple partitions for a Linux installation as you become more experienced and work on larger and more complex Linux systems:

- Placing the core operating system files on a separate partition allows you to upgrade the operating system without disturbing user data files or applications.
- Placing the core operating system files on a separate partition prevents user data files or applications from filling all hard disk space that the operating system needs in order to continue functioning.
- Storing user data on a separate partition may make it easier to create backups of that data.
- Separating data onto multiple partitions and setting up different options for how each partition is accessed allows you to implement certain security features. (This is described in Chapter 8.)
- Having multiple hard disks working to retrieve data at the same time can improve performance.

Using `fdisk`

During the installation process you must prepare a target partition where Linux will be installed before the installation source files can be copied to that partition. The installation program often gives you the opportunity to start the `fdisk` utility to prepare a target partition. Using the `fdisk` utility within an installation program can be a little intimidating: the screen goes blank except for a single line with the text `Command (m for help):`, and if you make a mistake, you might erase everything on the hard disk.

When the installation program you are using starts `fdisk`, the utility reads the partition table from a hard disk. The **partition table** is the information that defines the size and location on the hard disk for each partition. You modify that partition information using commands within `fdisk`. Nothing is actually written to the hard disk until you tell `fdisk` to do so, at which point the partition information is updated on the hard disk.

Although `fdisk` doesn't provide much assistance to new users, you can always use the `m` command to list the available commands in `fdisk` and the `p` command to list the partitions currently defined for the hard disk. Figure 3-3 shows sample output for the `p` command.

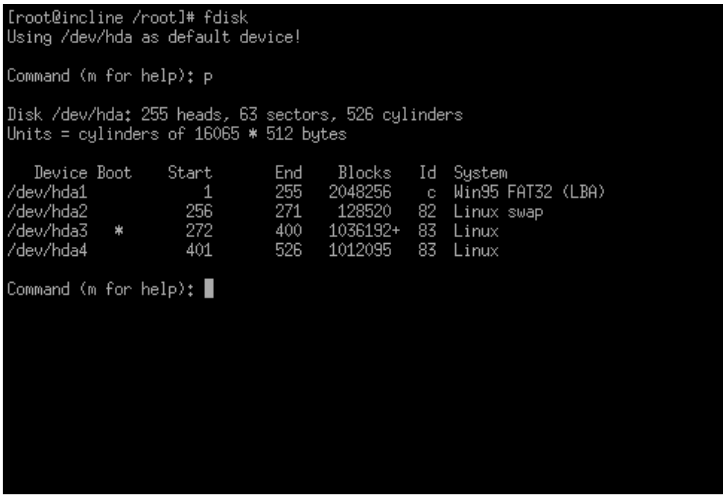

Figure 3-3 The `fdisk` utility showing a command list and a partition table

The steps for using `fdisk` to create a Linux partition are given below in two sets of steps. The first creates a Linux swap partition, the second creates a Linux native partition. Review these for now, and be prepared to follow them if necessary later, when you complete the hands-on projects at the end of this chapter.

In these steps it is assumed that you have a single hard disk and want to place the swap partition and main Linux partition on that hard disk.

1. When you have indicated to the installation program that you need to work with partitions (you must explicitly select `fdisk` as well), the screen changes to a black background and the message `Command (m for help):` appears. Type **`p`** to list the partitions currently defined on the hard disk.
2. If you used the `FIPS` program in Chapter 2 to create two separate Windows partitions (the second one being empty), study the partition information listed by the `p` command to determine which partition contains Windows data and which is empty. Because the empty Windows partition is the space you will use for Linux, you must delete that empty Windows partition before creating a new Linux partition. Enter **`d`** to delete a partition. When prompted, enter the number of the partition you want to delete.

3. Enter `n` to begin creating a new partition.
4. Enter `p` for primary partition. (You can create more than four partitions on a hard disk using extended, or logical, partitions, but you shouldn't need to do this.)
5. Enter the partition number you are creating. For this value, enter the next free partition on the hard disk. (For example, if only one partition is defined, for Windows, then the next available partition number is 2.)
6. Begin to define the size of the new partition by entering its first cylinder as requested. The range of available cylinders on the hard disk is listed in parentheses, like this:

   ```
   First cylinder (201-526, default 201):
   ```

   Enter the first number in parentheses as the starting point for the new partition.
7. Enter a size for the ending point of the new partition. Enter `+128M` to create a swap partition of 128 MB. You can choose the appropriate size for your installation based on the amount of RAM on the system and the size of the hard disk.
8. Enter `t` to set the type of the new partition.
9. When prompted, enter the partition number of the new partition you just created.
10. Enter `82` as the partition type when prompted for it. This is the code number for the Linux swap partition.
11. Use the `p` command again to see how the new partition is defined.

The size you enter for the swap partition is rounded to the size of the nearest cylinder, so it may not be exactly the size that you specified.

Next, you will need to create a native Linux partition (on which the operating system will be stored). Once again, review the following steps, and be prepared to use them later if necessary.

1. Enter `n` to begin creating a new partition.
2. Enter `p` for primary partition.
3. Enter the partition number you are creating by referring to the output of the `p` command and using the next available number.
4. When the size is requested, the range of available cylinders on the hard disk is listed in parentheses, like this:

   ```
   First cylinder (201-526, default 201):
   ```

   Enter the first number you see in parentheses. (For example, if the system displays the text above, enter the number 201.) This number specifies the starting point for the new partition.

You cannot use cylinder number 1 as the first cylinder of a partition because that cylinder is reserved for storing the partition table. If something seems to be wrong with the `fdisk` program because of an entry such as cylinder number 1, press Ctrl+C to exit `fdisk` without making any changes, and then run the program again.

5. Enter the second number in parentheses. This second number that you enter specifies the ending point of the new partition. To use the rest of the hard disk space for Linux, use the second number (526 in the sample output). If you want to create a Linux partition of a specific size, enter that size now (just as you did earlier for the swap partition). For example, to create a 700 MB partition, enter **`+700M`**.
6. Next, you need to specify the type of the new partition. To specify the partition type, enter **`t`**, enter the partition number for the partition you just created, and then enter **`83`**. (83 is the code number for a Linux native partition.)
7. Enter **`b`** to define the bootable or active partition on the hard disk.
8. Enter the partition number of the Linux native partition that you just defined.
9. Use the **`p`** command to see how the new partitions are defined. You should see an asterisk (`*`) next to the Linux native partition, indicating that the partition is active.
10. If you are satisfied with the partition changes you have made, write them to the hard disk by entering **`w`** for write. This command also exits the `fdisk` utility. If you used the numbers described in the steps, as prompted by `fdisk`, the remaining space on the hard disk is used for a newly created Linux partition when you enter the `w` command.

You can also exit the `fdisk` utility without making any changes to your hard disk by entering `q` for quit.

You may be prompted by the installation program to restart the computer after changing the partition table with `fdisk`. By restarting the computer you ensure that the partition table is correctly read by the installation program.

Using Disk Druid

The Red Hat Linux installation program includes a tool called Disk Druid. You can use Disk Druid instead of `fdisk` to set up partitions on the hard disk. The installation program prompts you to use Disk Druid before the installation source can be copied to a target partition. The utility provides the same basic functionality as `fdisk`, but within an easier-to-use interface. In addition to creating partitions, Disk Druid defines the mount points for Linux subdirectories. If you have already used `fdisk` or a commercial partition management product to create the Linux partitions (swap and native Linux) needed for the installation, you only need to define which partition on the hard disk should act as a mount point for the installation source files. This is the same Linux partition that you created in `fdisk` or Disk Druid, but you must manually select the partition at this point in the installation process.

Figure 3-4 shows the Disk Druid utility. The bottom half of the screen lists hard disks on the system. (On most systems, only a single hard disk, `/dev/hda`, is listed.) The top half of the screen lists the partitions on the selected (or only) hard disk. For each partition, a column notes the type of file system assigned to that partition (such as Linux swap or Linux Native, `ext2`), the size, and the mount point (if applicable, for native Linux partitions).

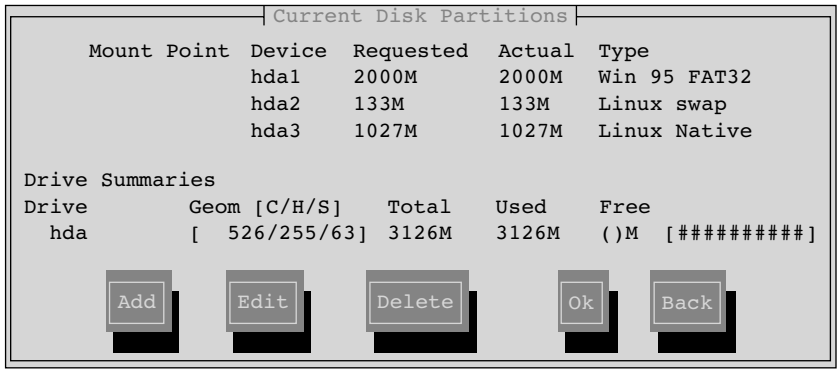

Figure 3-4 The Disk Druid utility in Red Hat Linux

Follow the next series of steps to set up new partitions in Disk Druid.

If you have not created a swap partition, create that partition first using these steps. The swap partition does not have a mount point because it is not accessed as part of the Linux directory structure.

1. Press the **Tab** key until the Add button is highlighted. The Add button begins the process of creating a new partition.
2. Press **Enter**. The Edit New Partition dialog box appears. Here, you can use the Tab and arrow keys to move to various fields.
3. Enter values for the mount point (this should be a forward slash, `/`), the file system type (this should be `ext2` to indicate a native Linux partition), and the partition size according to your preference.
4. If you want the partition size to use all remaining disk space, select the Grow to fill check box by pressing the Spacebar. (Don't use this option until you first create a swap partition.)

You may have created the necessary swap and native Linux partitions in another program such as `fdisk` or a commercial partitioning utility. In this case, you only need to follow these steps when Disk Druid starts:

1. Use the arrow keys to select the partition (in the top part of the Disk Druid window) that corresponds to the partition you want to use as the root of the Linux file system (the partition in which you want to install Linux).

2. Press **Tab** repeatedly until the Edit button is selected; then press **Enter**.
3. Enter `/` as the mount point for the partition.
4. Press **Tab** until the OK button is selected; then press **Enter**.

When you have created any necessary new partitions or, at a minimum, defined a partition that has `/` as a mount point, you can continue with the installation of Red Hat Linux.

Other Linux installation programs include their own specialized partitioning tools that are similar to Disk Druid. For example, Caldera OpenLinux provides a graphical tool that allows you to define partitions and mount points much as you do in Disk Druid. SuSE Linux uses the YAST configuration tool for setting up the target hard disk partition (and for selecting most other configuration options). Other versions of Linux use other tools.

Choosing What to Install

As you'll recall from Chapter 1, each Linux distribution comprises hundreds of different programs, including the Linux kernel, the Apache Web server, the GNU project system utilities, programming languages, compilers, graphical systems, games, and so forth. When installing Linux, you must decide which of these components to install, although the amount of freedom you have in making these choices varies by distribution.

Some distributions group the many possible components into a few different installation types. The **installation type** you choose determines which Linux software to install; the right installation type for your system depends on how the system will be used. For example, you might have a choice of installing a minimal system (which includes only the most basic components), installing a standard system (which includes the components the average user would be likely to use), or installing everything from the CD. Often, however, you will have more options regarding which software components you want to install.

Most Linux systems (including Red Hat, SuSE, Caldera, and Debian) gather many related files into a single software package. For example, all Linux products provide the Apache Web server, which is made up of dozens of files, as a software package. A **software package** is a single file that contains all the files needed to install and use an application or group of related applications. Special data formats are used to store many files in a software package. The Red Hat Package Manager format (abbreviated as `rpm`) is the most popular data storage format for creating software packages. Different data storage formats employed by Debian Linux also allow all of the files associated with an application to be stored in a single software package.

You will learn more about the `rpm` format in Chapter 4. Typical Linux products include between 400 and 1500 software packages. (To simplify matters, these packages are grouped into functional categories such as "text processing," "networking utilities," or "software development tool." Some Linux installations employ further generalizations based on a broad usage category for the Linux system. For example, a single selection such as "Web server installation" or "Desktop system" might define all of the categories, software packages, and files to install. During the installation you can specify which sets of packages or

which type of system you want to install. Different Linux distributions allow different levels of detail in this selection process. Figure 3-5 illustrates the variety of groupings, from general installation types to specific files.

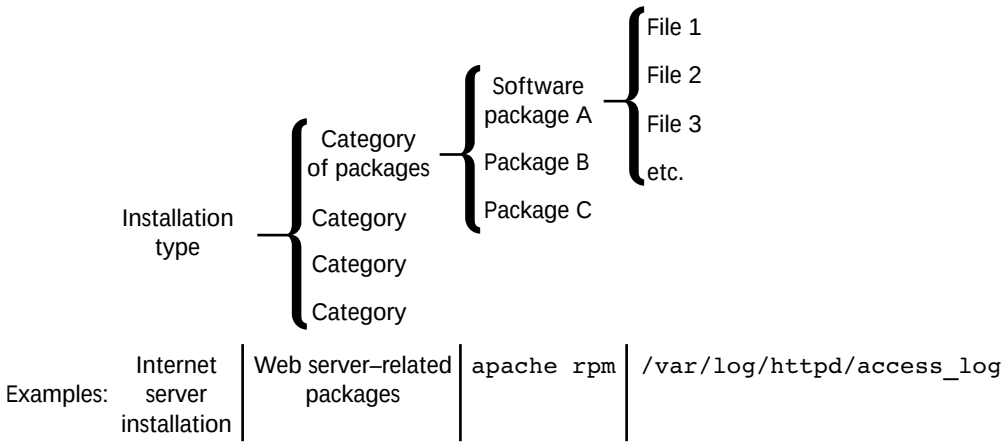

Figure 3-5 Files, software packages, categories of packages, and installation types

Red Hat Linux provides a few high-level installation types, such as Gnome Workstation, to create a desktop system using the Gnome graphical desktop. In addition, Red Hat Linux provides a Custom option that lets you select various categories of software packages according to your needs. The Custom option requires you to make decisions about 25 different categories of software packages. For each package, you need to decide whether to include or exclude it from the installation. Although it entails more work on your part, the Custom installation type is recommended for two reasons. First, it allows you to see the many options provided by the distribution you are installing. Second, it allows you to customize your installation to fit your needs. Table 3-4 shows the categories of packages (called Package Groups) that Red Hat Linux provides during a Custom Server installation.

Table 3-4 Package Groups in Red Hat Linux

Package Group	Description
Printer Support	Provides the ability to send documents to local or remote printers using standard Linux print server programs
X Window System	Provides the foundation for graphical applications in Linux. Required for the Gnome or KDE option
GNOME	Adds the Gnome graphical desktop interface to the system
KDE	Adds the KDE graphical desktop interface to the system
Mail/WWW/News Tools	Allows client access to the Internet for reading e-mail, browsing the Web, and reading newsgroups
DOS/Windows Connectivity	Includes programs that help you connect to DOS or Windows systems on the same network

Table 3-4 Package Groups in Red Hat Linux (continued)

Package Group	Description
Graphics Manipulation	Provides software for viewing and creating graphics files in various formats
Games	Adds a variety of Linux games to the system
Multimedia Support	Includes drivers and other software for using a sound card and playing video clips within Linux
Networked Workstation	Provides a collection of utilities considered useful if you are using Linux on a local network (such as an Ethernet or Token Ring network)
Dialup Workstation	Provides a collection of utilities considered useful if you are using Linux as a client that dials a modem to connect to a network
News Server	Provides Usenet newsgroup server capabilities
NFS Server	Provides the ability to make local file systems (hard disk partitions) available to users working on other computers
SMB (Samba) Server	Provides the ability to make local Linux resources available across the network to users working on Windows-based computers
IPX/NetWare Connectivity	Provides the ability to connect to NetWare servers located on the same network as the Linux system
Anonymous FTP Server	Lets Linux act as an FTP server to provide file downloads to networked users
Web Server	Installs and activates an Apache Web server on Linux
DNS Name Server	Allows Linux to act as a DNS server for other computers on the same network that need to convert domain names to IP addresses
Postgres (SQL) Server	Provides a complete client/server database on Linux (using the freely available Postgres program)
Network Management Workstation	Installs utilities considered useful for someone using the Linux system to manage other systems on a local network
TeX Document Formatting	Provides text-processing capabilities using the TeX document language (commonly used on UNIX systems)
Emacs	Installs the large and powerful Emacs text editor
Development	Provides software components needed to develop Linux software or compile Linux source code files
Kernel Development	Provides software components needed to change the Linux kernel after installation
Extra Documentation	Installs additional documentation files
Utilities	Installs additional useful Linux command-line and graphical utilities
Everything	Installs everything on the Red Hat Linux CD-ROM

Red Hat Linux also allows you to select individual packages to install, but choosing among 600 or so packages takes more effort than most people can invest during the installation process. You'll find it easier to install package groups when installing Linux. Then you can easily add or remove individual software packages after completing the installation.

User Accounts

The Linux installation process creates an administrative user account named `root`, which is used for system management. Whoever has access to this account can control the entire system. During installation, you must specify a password for the `root` user account. Choose this password carefully and guard against anyone discovering the password. Because the `root` account is so powerful, you should only use it to complete system administration work.

For security purposes, you will not see the `root` password on screen as you type it. This prevents anyone from observing the password as you type. You will be prompted to enter the password a second time to make sure you typed it correctly.

The `root` account is created automatically—you only have to provide the password. In some distributions (including Red Hat Linux) you may also be prompted to create another user account besides `root`. If you are prompted to do so, choose a brief username of eight characters or less (such as nwells, thomasj, jane, or rms). Depending on your version of Linux, you may be prompted to enter a full name (your own complete name, such as Nicholas Wells), a password, or other information. You can use the commands described in Chapter 8 to change or add information to this user account after you have completed the installation.

Configuring the Graphical System

As mentioned in Chapter 2, configuring the graphical system (the video card) is the most challenging part of most Linux installations. The Linux installation program is designed to configure the video card for you, but you may need to enter some of the hardware information you gathered in Chapter 2 in order to complete the installation. If you are using the standard XFree86 graphical software included with all Linux distributions, you'll find the configuration file located at `/etc/X11/XF86Config` on a Red Hat Linux system, or `/etc/XF86Config` on some other Linux systems.

If the graphical configuration that you set up during the installation process does not appear to work correctly after starting the new Linux system, try entering one of the following commands (at a command line) to reconfigure the graphical system. (Keep in mind that not all of these commands are available on all Linux systems.) Note the upper- and lowercase letters used in each command name.

- `Xconfigurator`
- `lizardx`
- `XF86Setup`
- `xf86config`

Configuring the Boot Loader

During the Linux installation you will have the option of installing the Linux boot manager (LILO) in one of several locations. The option that you select depends on how you have configured the other operating systems on the computer and on how you want to start Linux. If you are uncertain, you can select the option that is preselected by the installation program. This option is based on the status of your system as determined by the installation program when it probes your hard disk. The following list explains reasons for choosing each of the possible locations for LILO:

- The Master Boot Record or MBR: this location ensures that Linux boots correctly after a new installation. When LILO is installed in the MBR, control passes from the BIOS of the computer directly to LILO, which can then start Linux directly. Many system administrators try to avoid writing LILO to the MBR because it may interfere with other operating systems used on the same computer, such as Windows NT.
- The boot sector of the partition on which Linux is being installed: this is the recommended location for LILO. When LILO is installed here, control passes from the BIOS to the MBR, to the copy of LILO stored on the boot sector of the active partition. LILO can then start Linux or pass control to another partition to start another operating system, such as Windows 2000. This method assumes that the partition on which Linux was installed is marked as the bootable or active partition. If it is not, control will pass from the BIOS to the MBR to the active (non-Linux) partition, and LILO will never be invoked (hence Linux cannot be started). If Linux is not installed on the active partition, the MBR or the active partition must contain a boot manager of some type (BootMagic, another copy of LILO, or another boot manager program) in order to have the option of starting Linux on the nonactive partition that you are installing.
- A floppy disk: this allows you to keep the MBR intact and leave another partition (such as one containing Windows) as the active partition while allowing you to start Linux from a floppy disk. When LILO is installed on a floppy disk, you can only start Linux by inserting that floppy disk into the computer. Then when you turn on the computer, control passes from the BIOS to the floppy disk containing LILO. LILO then starts Linux (or other operating systems if so configured). If the floppy disk is not inserted, the Linux operating system cannot be started and has no effect on other operating systems stored on the computer.

When you start a computer system, you see the word `LILO` appear very briefly on the screen as the LILO program is loaded and executed. Then a prompt appears at which you can enter the name of the operating system you want to start. The prompt used by LILO is `boot:`. When you see this prompt, you can press the Tab key to see a list of the operating system names that LILO recognizes. This list includes `linux`. It also includes either `dos` or `win` if you have a computer system on which Windows is also installed.

STARTING LINUX

After you have installed Linux, you can reboot the computer and begin using the operating system. The following sections describe how to begin working with the newly installed operating system.

Logging In

You cannot do anything on a Linux system until you have logged in using a valid username and password. The term **logging in** refers to the process of identifying yourself as a valid user who has been assigned a certain set of access rights. **Authentication** is another, more precise, term for logging in. Authentication also refers to entering a valid username and password, but it implies that the system has verified the identity of a user based on specific rules.

Although you should not use the administrative account, named `root`, for normal work, it may be the only account available on your system immediately after installation. In this case, use the username `root` and the password you entered during the installation to log in to your system, and immediately create another user account as described later in this chapter.

Depending on the configuration of your new Linux system, you might see a character-based login screen like the one shown in Figure 3-6, or a graphical login screen like the one shown in Figure 3-7. The character-based login screen starts a command-line interface after you log in; the graphical login screen starts a graphical desktop interface after you log in.

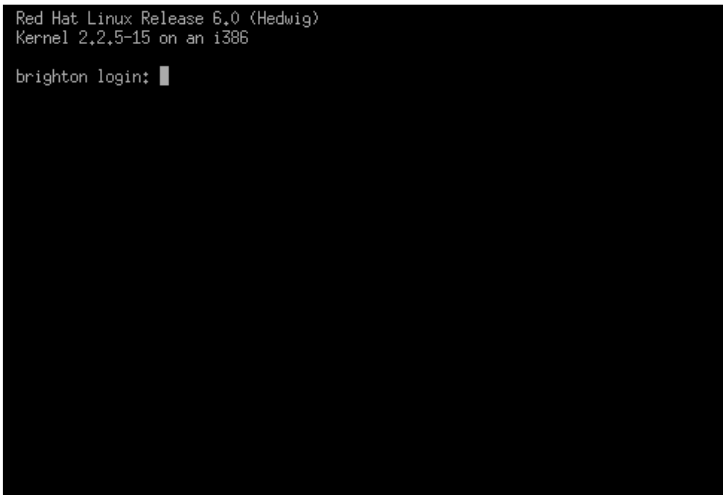

Figure 3-6 Character-based login screen

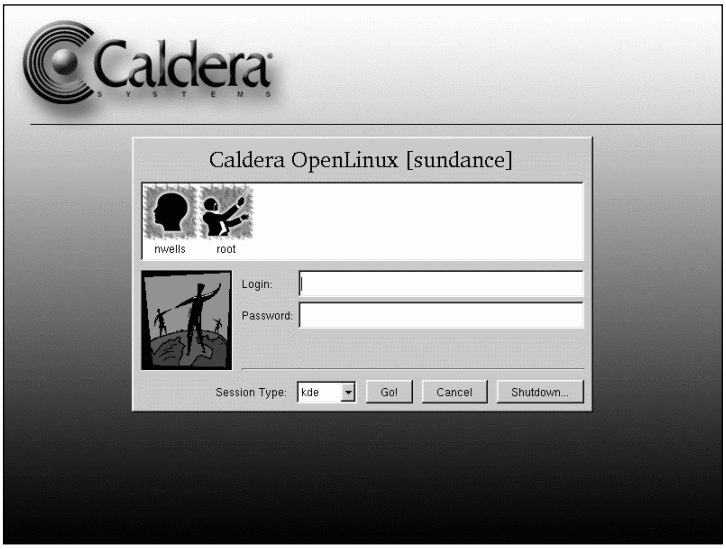

Figure 3-7 Graphical login screen

Starting the Graphical System

Linux systems that use a graphical login screen similar to the one shown in Figure 3-7 will change immediately to a graphical display after you enter a username and password. After logging in at a character-based screen, you can enter a command to start the graphical system. On most Linux systems you can use the following command (in all lowercase letters) to start the graphical system:

```
startx
```

Of course, in order to use this command you must have previously configured the graphical system. This is normally done during the installation process, although it is a challenging part of the installation and may require additional work after completing the installation.

Most recent Linux distributions provide a desktop graphical interface known as KDE. The desktop provided by Red Hat Linux is called Gnome. Figure 3-8 shows the Gnome desktop as it would appear after a new installation. Chapter 5 describes the graphical system in detail and explains how to configure it.

Figure 3-8 The Gnome desktop interface

Creating the First User Account

Many Linux distributions (including Red Hat Linux) create a regular user account during installation. A **regular user account** is a user account that is not used for system administration work (as the `root` account is). A regular user account has a name similar to a person's name, such as nwells, or georgew.

You should always use a regular user account to log in to Linux. If you did not define a regular user account during the installation, create one immediately after you log in to the new system as `root`. After creating a regular user account, use it for all work except system administration tasks. Note that you can only create a new user account when you are logged in as the `root` user.

Any one of several utilities will allow you to create a new user account. The most common command-line utility for creating user accounts is called either `useradd`, or `adduser` (depending on your system). On Red Hat Linux, you can create a new user account by combining the `useradd` command with a new user account name. For example, to create a new user account named `nwells`, you would log in as `root`, and then enter the following command:

```
useradd nwells
```

Managing user accounts is a large part of the work of a system administrator. Chapter 8 describes in detail how to set up user accounts with many different optional settings and how to manage those accounts on a busy Linux system.

Testing Network Connections

Networking is configured as part of the installation process. Many users will want to begin using the network connection immediately after installation, in order to send e-mail, browse the Web, or otherwise communicate with other networked users. Before using the connection, though, try a few simple commands to make sure it is correctly configured. Note that you do not have to be logged in as `root` to use the commands discussed in this section.

The **`ping`** command is used to test a network connection. To test your connection, begin by entering the following:

```
ping 127.0.0.1
```

You should see lines appear on screen once every second. Press Ctrl+C to stop the command. If no lines appear, networking is not configured or activated on your system. If the first `ping` command worked, try this next command, using the IP address that you entered for your system during the installation. (If you did not enter an IP address during installation, networking is not yet configured, and these tests do not apply.)

```
ping <IP address>
```

If you see lines appear once every second, your network adapter card has been configured. If no lines appear, see the following troubleshooting sections for suggestions on configuring your network adapter. Next, try to contact a remote site using an IP address. (You can use any IP address that corresponds to a real server on your local network or on the Internet, if you have an Internet connection.) For example, you might try the following command:

```
ping 207.49.12.1
```

If lines appear once every second, you have established a connection with another computer. Finally, try using a domain name with the `ping` command:

```
ping www.yahoo.com
```

If lines appear once every second, you are ready to use the network. If this last test fails (that is, if no lines appear), you need to configure your DNS server address. See the following troubleshooting sections.

To access the Internet using a Web browser, you may also need to configure your Web browser to use different server names or addresses for security purposes. Ask your system administrator or ISP if any configuration changes are required.

The preceding information outlines only the most basic network information. Complete details on configuring and using networking are beyond the scope of this book.

TROUBLESHOOTING A NEW INSTALLATION

Depending on your Linux distribution and your computer hardware, the installation process may go very smoothly, or you may have to overcome numerous obstacles. The following sections provide guidance on how to solve some common problems you might encounter during a difficult Linux installation. Not all possible problems are covered here, of course. In fact, some problems may not have solutions, because a system may have hardware that is not supported by Linux. (You would normally discover this while you gathered information about the computer before installation, as described in Chapter 2.) In addition, some troubleshooting requires knowledge far beyond what has been presented so far in this book. If your efforts don't succeed, ask a Linux expert for help.

The System Won't Boot

If the Linux installation program will not start, the problem may lie with the boot disk, which may not contain all the necessary files. You might also be trying to start the installation from a device that is not used by the BIOS during the system start-up. For example, if the BIOS is configured not to look at the floppy drive before booting from the hard disk, using a boot disk won't start the installation program. In this situation you must alter the BIOS configuration so that it checks the floppy drive before passing control to the hard disk.

Once you install Linux, you may have problems getting the newly installed Linux system to boot. (Or perhaps the computer will not boot at all.) In this situation, use the boot disk you created during the installation process to start the computer. The Red Hat Linux boot disk will search your hard disks for an existing Red Hat Linux installation and start that operating system after a few seconds.

After installing Linux and starting the system from a boot disk, you may also need to specify the partition on which the operating system is stored, via a boot parameter. A **boot parameter** is a piece of information that you can type in at the LILO prompt. The information you enter is sent to the Linux kernel as the system is being booted. These parameters are normally used to affect how Linux recognizes hardware devices, or which features of the operating system are enabled.

Each operating system that LILO can start has a label (a name) associated with it. When the LILO prompt (`boot:`) appears after you start the system, you can enter a label to start the corresponding operating system. For example, enter the label `linux` to start Linux. At this point you can also add a boot parameter to control the Linux kernel. One example of a boot parameter is `root=`, which designates the root partition for the operating system being booted. The `root=` parameter requires that you specify the partition to use as the root partition for Linux. This example shows what you could enter at the `boot:` prompt to use `/dev/hda1` as the root partition:

```
linux root=/dev/hda1
```

If you see only part of the word `LILO` (such as just the letters `LI`) as the computer system boots, and nothing else happens, the boot manager has been incorrectly installed (this is a rare occurrence). In this case, use the steps described in Chapter 4 to update the configuration of LILO after accessing the new Linux system from a boot disk.

The Graphical Interface Doesn't Work

Setting up the graphical system of Linux is often challenging—so challenging that an entire chapter of this book (Chapter 5) is devoted to the task. Before you can delve into the details of setting up the graphical system, you need to make sure your distribution supports the video card you want to use. Because virtually all Linux distributions use the XFree86 software to support the X Window System, you can check the version of XFree86 software included with your copy of Linux (ask the vendor if you don't have product literature to review), and then check the Web site *www.xfree86.org* to see if your video card is supported by that version of XFree86.

Some video cards are not supported by XFree86, but many of these cards are supported by commercial X Window System products that you can purchase and add to your Linux system. These products are available from Xi Graphics (see *www.xig.com*) and MetroLink (see *www.metrolink.com*).

Chapter 5 provides more explanation about how the graphical system works and how to configure the X Window System. You can always finish the installation of Linux without configuring graphics, then set them up afterwards. Even if your video card is listed as supported by XFree86, you may have trouble getting it correctly configured. In these cases, you may be able to find a Linux expert in your school or in a Linux user group who has experience with the video card you are trying to configure.

Some Hardware Isn't Available

Sometimes hardware that is included on the list of supported hardware on the Linux vendor's Web site is nevertheless not available after starting Linux. This is because the hardware is not correctly configured. As with a system that won't boot, the configuration can generally be corrected by adding a boot parameter when starting Linux. For example, some computers running Linux will not access the CD-ROM drive correctly unless the device name is added as a boot parameter, like this:

```
linux cdrom=/dev/hdc
```

Or Linux may not access all of the available system RAM because of limitations in the computer's BIOS. In this case, you can tell Linux the amount of RAM on the system using this format:

```
linux mem=<<amount of memory>>M
```

So to indicate 128 MB of RAM, you would use the following command:

```
linux mem=128M
```

Be sure to use the correct value for the amount of RAM on your computer, or Linux will crash as it tries to work with nonexistent memory.

You can also combine multiple boot parameters on a single line separated by spaces. For example:

```
linux cdrom=/dev/hdc mem=128M
```

The Boot Parameters HOWTO document provides details about how to add these parameters to make your hardware function correctly in Linux. (Chapter 1 describes how to access HOWTO documents on the Red Hat Linux CD or via the Internet.)

CHAPTER SUMMARY

- Different Linux distributions vary in how the Linux files are arranged, which installation programs are provided, and a few other areas; but at their core, all Linux distributions use the same set of software taken from the Internet. The installation programs used by different Linux vendors are also quite similar in the information that they request during the installation process.
- The process of installing Linux involves starting an installation program (normally from floppy disk or CD-ROM) and then answering questions about where Linux should be installed and how it should be configured. During the installation you prepare hard disk partitions for Linux using a tool such as `fdisk` or the Red Hat Disk Druid utility, assign a password to the `root` user account, and decide which types of software packages to install.
- After installing a new Linux system, you must log in using a valid user account name. You can then start the graphical system, check networking, and create an additional user account if necessary.
- Troubleshooting a Linux installation that is not going well may involve adding boot parameters, researching additional hardware information, or using configuration tools to finish setting up features such as the graphical environment or network access after the installation is otherwise complete.

KEY TERMS

authentication — The process of identifying a user to a computer system via some type of login procedure.

boot disk — A floppy disk containing a portion of the Linux installation program, which can be used to start the Linux installation program.

boot parameter — A piece of information passed directly to the Linux kernel as the system is being booted. These parameters are normally used to affect how Linux recognizes hardware devices or to enable certain features of the operating system.

bootable CD-ROM drive — A CD-ROM drive that can launch an operating system (or other program) directly from a CD without accessing the hard disk. (This feature of the CD-ROM drive must be enabled by the BIOS.)

disk image — A single file that contains an exact copy of a floppy disk.

`fdisk` — A utility used to create hard disk partitions and configure how they are used.

install disk — A disk used to start the Linux installation program on some distributions of Linux. *See* boot disk.

installation source — The set of files from which Linux is installed. These files are normally stored on a Linux CD.

installation type — A specification indicating which Linux software to install; the correct installation type depends on how the Linux system will be used.

logging in — The process of identifying yourself as a valid user who has been assigned a certain set of access rights.

mount point — A subdirectory through which a set of data is accessed.

partition table — Information that defines the size and location of each partition on a hard disk.

`ping` — A command used to test a network connection.

regular user account — A user account that, unlike the `root` account, is not used for system administration work. A regular user account has a name similar to a person's name.

software package — A single file that contains all the files needed to install and use an application or group of related applications. Special data formats are used to store many files in a single software package.

target hard disk partition — The location on the system's hard disk where Linux will be installed. Also known as *target partition*.

Review Questions

1. Linux distributions vary greatly in the core functionality that they provide. True or False?
2. The installation program included with a Linux distribution is usually created by:
   a. The company, or vendor, that sells the Linux distribution
   b. A group of many vendors working together to create a common installation program
   c. The Gnu project of the Free Software Foundation
   d. Linus Torvalds, as part of the Linux kernel itself
3. Name two technical differences and two nontechnical (marketing related) differences between various Linux products.
4. Explain why installing a new Linux system is different from installing an application such as a spreadsheet or a database package.
5. The target hard disk partition is where:
   a. A dual-boot Windows system resides
   b. The Linux operating system will be installed
   c. Back-up data must be stored for Linux to access it
   d. The Linux installation program is stored
6. Possible locations for the installation source data do *not* include which of the following:
   a. The target hard disk partition
   b. A local CD-ROM
   c. A local hard disk
   d. A networked server using the SMB protocol

7. To start a networked installation you should contact:
   a. Red Hat software to obtain a different boot disk
   b. Your system administrator to obtain the target partition for the installation
   c. Your system administrator to determine whether your hard disk supports a network-based installation
   d. The network administrator of the server containing the installation source
8. When you turn on a computer, _____ sends control to the MBR of the first hard disk, or to another device such as a bootable CD-ROM drive or a floppy drive.
   a. Linux
   b. the `fdisk` utility
   c. LILO
   d. the BIOS
9. A boot disk used to start a Linux installation program may have different names on different distributions of Linux. True or False?
10. In which circumstance would you need a boot disk to start the Linux installation program?
    a. The computer does not have a bootable CD-ROM drive.
    b. When installing on a laptop.
    c. The root disk has become corrupted.
    d. When installing from a network installation source.
11. Name the tool used in Windows to copy a disk image to a floppy disk.
12. The `fdisk` utility is used to:
    a. Prepare partitions on a hard disk
    b. Create a boot disk from a disk image
    c. Start a Linux installation program from a boot disk
    d. Launch Linux from a Windows-based fixed disk (hard disk)
13. The device name `/dev/hda3` would refer to:
    a. The third partition on the first IDE hard disk
    b. The third partition on the first SCSI hard disk
    c. The third IDE hard disk
    d. The swap partition stored on a boot disk
14. Name three reasons why you might place different parts of the Linux file system on different hard disk partitions.
15. The _____ defines the size and location on the hard disk of each partition.
    a. MBR
    b. `fdisk` utility
    c. partition table
    d. Disk Druid utility

16. Most Linux distributions use a system of software packages to make managing software easier and more efficient. True or False?
17. Explain why you might choose to select groups of software packages rather than individual software packages during an installation.
18. The _____ user account is created as part of every Linux installation process.
    a. LILO
    b. installation source
    c. `useradd`
    d. `root`
19. Storing the LILO boot manager on a floppy disk allows you to:
    a. Start Linux only when the floppy disk is inserted in the computer
    b. Slow down the boot process to examine how it works
    c. Store user accounts along with LILO on the boot floppy
    d. Install Linux in a more secure manner
20. The _____ command normally starts the graphical system when working in a text-based Linux system.
    a. `startx`
    b. `GNOME`
    c. `KDE`
    d. `gnome`
21. The `ping` command is used to:
    a. Test a network connection
    b. Test the sound card configuration
    c. Test the integrity of a hard disk connector
    d. Send a small e-mail message to another system
22. Boot parameters can be used to pass additional information to the Linux kernel as the Linux installation program is started. True or False?
23. Name three programs that can be used to configure the X Window System after the Linux installation is completed.
24. Which of the following is *not* likely to cause hardware problems in Linux?
    a. The correct module supporting that hardware is not yet installed.
    b. A parameter identifying the hardware was not entered correctly.
    c. Windows has disabled the hardware for use under Linux.
    d. The hardware is not supported by Linux.
25. The surest way to have LILO start Linux is to install LILO on:
    a. The MBR
    b. The boot sector of the active partition
    c. A floppy disk
    d. The installation source partition

HANDS-ON PROJECTS

Project 3-1

In this activity you create a boot disk to use when installing Red Hat Linux. To complete this activity you should have a computer with Windows 95 or 98 installed and both a CD-ROM drive and a floppy disk drive. You should also have a Red Hat Linux CD like the one included in this book and a blank floppy disk. For this project, it is assumed that you will install Red Hat Linux from a local CD-ROM drive rather than from a network connection and that you are not using a laptop.

1. Start Windows.
2. Insert the Red Hat Linux CD in the CD-ROM drive and a blank floppy disk in the floppy disk drive.
3. Double-click the **My Computer** icon on the Windows Desktop. A window opens showing you the contents of your computer.
4. Double-click the icon for the CD-ROM drive within the My Computer window. A window appears showing you the contents of the Red Hat Linux CD.
5. Double-click the **dosutils** folder within the CD-ROM drive window. A window appears showing you the contents of the Images folder.
6. Double-click the **rawrite** program within the Images folder window. An MS-DOS window appears in which the rawrite program is executed. The following text line is displayed:

   ```
   Enter disk image source file name:
   ```

7. Enter the name of the disk image that you are copying to a floppy disk. The exact text you enter depends on the drive letter for your CD-ROM drive in Windows. If your drive letter is E:, for example, enter this text: **`E:\images\boot.img`**. The following text line is displayed:

   ```
   Enter target diskette drive:
   ```

8. Enter the drive letter for the floppy disk drive on the Windows system. This is normally the letter *A*. The following text line is displayed. (The value `A:` varies based on the drive letter that you enter in this step.)

   ```
   Please insert a formatted diskette into drive A: and press
   ENTER- :
   ```

9. Check that the floppy disk is inserted, and then press **Enter**.
10. When the disk image has been copied from the CD to the floppy disk, the floppy disk drive stops turning (and the light on the front of the drive will turn off). Eject the floppy disk and close the MS-DOS window on the Windows Desktop. Leave the Red Hat CD in the CD-ROM drive.
11. Shut down the Windows computer.

Project 3-2

In this activity you install Red Hat Linux 6.0 from the CD included with this book. To complete this activity you should have the boot disk you created in Project 3-1, the Red Hat Linux CD included with this book, a second blank floppy disk to use during the installation, and a computer with sufficient resources to complete the installation (about 700 MB of hard disk space and 16 MB of RAM). To perform the steps that follow you must have sufficient free space on your hard disk so that you can use the Disk Druid tool to create a swap partition and a native Linux partition during the installation.

Installing Linux is a lengthy procedure, and there may be slight differences in some of the following steps because of differences in the hardware installed on your system. The steps given here along with the on-screen menus should guide you through the entire process. If your system has a SCSI adapter card, the dialog boxes and order of the steps may vary slightly on your computer. Because few standard PCs include SCSI adapter cards, these steps reflect an installation on an IDE hard disk.

1. Verify that the Red Hat Linux CD is still in the CD-ROM drive and that the boot disk that you created in Project 3-1 is inserted into the floppy drive, and then turn on the computer. After a few moments, several lines of text appear on a black background with the title `Welcome to Red Hat Linux` at the top of the screen.
2. Press **Enter** to continue. Messages display on the screen as different parts of the Linux installation program are loaded. (This loading process may take as much as three full minutes, depending on your computer's speed.)
3. When the installation program has finished loading, you see a blue screen with a welcome message. Review the message and press **Enter** to continue.
4. A list of language options appears, as shown in Figure 3-9. This same type of selection screen is used throughout the Red Hat Linux installation program.

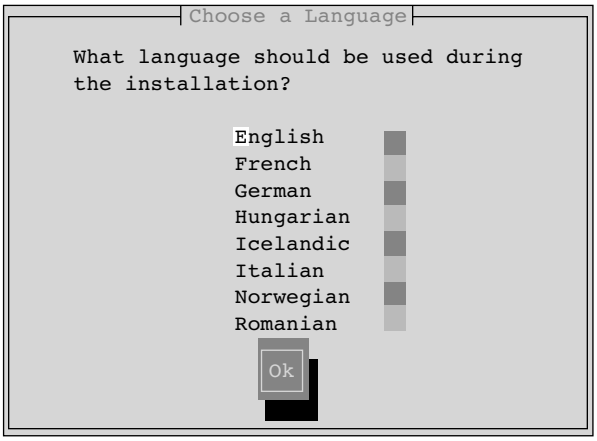

Figure 3-9 The `Language Selection` screen in the Red Hat Linux installation

5. Use the Up and Down arrow keys to select a language for the installation. (These steps are based on the English version.) Press **Enter** to continue. The `Keyboard Type` dialog box appears.

6. Use the Up and Down arrow keys to select the keyboard you are using from the list shown. (The items in the list are quite cryptic—the selection `us` is appropriate for most keyboards in the United States.) Press **Enter** to continue.
7. If PCMCIA laptop hardware is detected on your computer, a dialog box labeled `PCMCIA Support` appears. If you were installing Red Hat Linux on a laptop and required a connection via a PCMCIA expansion card to a network or SCSI adapter, you would need to choose `Yes` and use a special PCMCIA disk (created from the Red Hat Linux CD as in Project 3-1). Choose `No` to continue.

Notice that a `Back` button appears at the bottom of most dialog boxes. You can return to a previous dialog box by pressing Tab repeatedly until the `Back` button is highlighted in white; then press Enter. To advance to the next screen, you can either press the F12 key or press Tab until the `OK` button is selected, and then press Enter.

8. The `Installation Method` dialog box appears. In this dialog box you can specify where the Red Hat Linux data files are located. Use the arrow keys to select `Local CDROM`, and then press **Enter** to continue.
9. A message box tells you to insert your Red Hat Linux CD into the CD-ROM drive. (If you followed Step 1 of this project, the CD should already be in the CD-ROM drive.) Press **Enter** to continue. A message tells you that the CD-ROM is being initialized.
10. The `Installation Path` dialog box appears. In this dialog box you select whether you want to install a new system or upgrade an existing Red Hat Linux system. The `Install` button is selected by default. Press **Enter** to continue.
11. The `Installation Class` dialog box appears. Use the arrow keys to select `Custom`. (Other options include `Workstation` and `Server`. The `Custom` option allows you flexibility later regarding which components you want to install.) Press **Enter** to continue. Depending on your system, the `SCSI Configuration` dialog box may appear next.

You may see a different series of dialog boxes related to the process of configuring a SCSI card than described here.

12. If your system had a SCSI card, you would choose Yes and select the appropriate SCSI configuration from a list. No is selected by default because very few standard PCs have SCSI adapters. Make the correct selection, and press **Enter** to continue. The `Disk Setup` dialog box appears.
13. Make certain that the `Disk Druid` button is selected. (Use the Tab key if necessary to select it.) Press **Enter** to continue. The Disk Druid hard disk management utility appears, similar to the screen shown in Figure 3-4.

14. Review the Type column on the right side of the screen. If a Linux swap partition and a Linux native partition are not shown, press **Tab** until `Add` is highlighted in white, and then press **Enter** to begin creating the needed partitions for the installation. (These steps assume that the partitions have already been created using Disk Druid, `fdisk`, or another tool, but at this point in the installation process you can use the steps outlined in the chapter text to create new partitions on the hard disk.)
15. In the top part of the window, use the Up and Down arrow keys to select the partition on which you want to install Red Hat Linux. This partition must be labeled as "Linux native" in the Type column.
16. Press **Tab** repeatedly until the `Edit` button is highlighted in white.
17. Press **Enter** to open the `Edit Partition` dialog box.
18. In the Mount Point field, enter `/` (a single slash character) and press **Enter**. The `Edit Partition` dialog box closes, and you see the Disk Druid main screen again.
19. Press **F12** to continue the installation process. The `Active Swap Space` dialog box appears.
20. Make certain that the partition number you defined in Disk Druid as Linux swap space is marked with an asterisk, so that it can be formatted. Press the Spacebar if an asterisk does not appear to the left of the device name. (You don't need to select the `Check for bad blocks during format option`.) Press **Tab** until `OK` is selected, and then press **Enter** to continue. The `Partitions To Format` dialog box appears.
21. If you had selected multiple Linux native partitions with different mount points by using the Disk Druid tool in Steps 14 through 18, each one would be listed here. Here it is assumed you are using one partition for the entire installation, so a single partition (for the / directory) is listed in this dialog box. Make certain that the box to the left of the partition name contains an asterisk. (If it does not, press the Tab key repeatedly to select that line, and then press the Spacebar to add an asterisk.) You do not need to place an asterisk in front of the `Check for bad blocks during format` option unless you are concerned about the integrity of the hard disk on which you are installing Linux. Using this option adds significantly to the time required to prepare the partition for installation. Press **F12** to continue. The `Components to Install` dialog box appears.
22. Using the information provided in the chapter text and your own preferences, select the items you want to install by moving up and down the list with the arrow keys. Press the **Spacebar** to select or unselect each item as you highlight it. If you have sufficient hard disk space (more than 1 GB on the Linux partition), install as much as you can so you have more tools available to experiment with and learn about Linux. Don't select the option labeled `Select individual packages`. When you have finished, press **F12** to continue. The `Install log` dialog box appears.
23. Review the information in the `Install log` dialog box, which tells you that you can see information about the installation after it is completed by looking at the file `/tmp/install.log`. Press **F12** to continue. Message boxes inform you that the file system is being formatted and that packages are being installed. After a moment, the `Package Installation` screen appears. You don't need to do anything at this point.

24. Watch the packages being installed within the `Package Installation` screen. Note that the package names and sizes are shown as each one is installed. The total number of packages to be installed (based on your selection in Step 22) and the total size of all packages is shown in the middle of the screen. Read a few of the `Summary` lines as packages are installed. This part of the installation takes from 5 to 30 minutes depending on the speed of the computer system.
25. After the package installation is completed, a message box informs you of any mouse device that the installation was able to locate. Press **Enter** to continue. The `Configure Mouse` dialog box appears.
26. Select the type of mouse you have installed from the list shown by using the arrow keys. Press **Tab** to select the `Emulate 3 Buttons` item. Press the **Spacebar** to select that item. Then press **F12** to continue. The `Network Configuration` dialog box appears.
27. If you want to configure an Ethernet or other network adapter card, make certain that the **Yes** button is selected, and press **F12** to continue. (Most Linux systems are connected to networks, but you can choose to configure networking after the installation if you prefer.) A `Load module` dialog bo appears. If you prefer to choose No and configure networking after the installation is completed, skip to Step 31.

If the installation program detects a network card, you may not see some of the dialog boxes described here, or you may see different dialog boxes than described here.

28. In the `Load module` dialog box, use the arrow keys to select the module that corresponds to your network adapter card. If the installation program can detect a certain network card, the module for that card is selected by default. Press **Enter** when you have highlighted a module. If the module you select is not the correct one, you can return to this screen and try another module.
29. In the `Module options` dialog box, leave the `Autoprobe` option selected so that the installation program will attempt to autoconfigure the network card. You can also choose `Specify Options` and enter IRQ and I/O port numbers if the `Autoprobe` option is not successful.

You can install Red Hat Linux on a laptop using the CD included with this book. But because PCMCIA support is not part of the standard installation program, you will not be able to configure networking until you have completed the installation and rebooted the new Linux system. If you need to access the PCMCIA devices on the laptop in order to complete the installation via a network or SCSI connection, you must create a `bootnet` or `pcmcia` disk from the image files on the Red Hat Linux CD. See the documentation on the CD for more information.

30. After successfully installing the network adapter module, follow the on-screen instructions to enter the networking parameters specific to your local network.

31. The `Configure Timezones` dialog box appears. Use the arrow keys to select your time zone from the list shown. If your computer is only running Linux (and is not sharing a hard disk with Windows), select the option `Hardware clock set to GMT` by using the Tab key and pressing the Spacebar. Press **F12** to continue. The `Services` dialog box appears.
32. Use the arrow keys to highlight each service that you want to select or unselect, pressing the **Spacebar** to add or remove the asterisk next to each item in the list. In most cases, you can leave the services list configured as it is by default. Press **F12** to continue. The `Configure Printer` dialog box appears.
33. Use the **Tab** key to select **No**. (Printer configuration is discussed at length in Chapter 13.) Press **Enter** in the `Configure Printer` dialog box to continue. The `Root Password` dialog box appears.
34. Enter a root password in the `Password` field. Nothing appears as you type, though a flashing cursor is shown. Press **Tab** to return to the `Password` field, and reenter the password. Press **F12** to continue. You cannot continue until you have entered two matching passwords.
35. The `Authentication Configuration` dialog box appears. The default settings in this dialog box define how Red Hat Linux will store password information to protect it from intruders. Because the default settings provide good security, you can accept them by pressing **F12** to continue. (The `Enable NIS` check box is used if you are located on a network that uses the network information system protocol. Ask your network administrator if you think this might apply to the system you are installing.)
36. The `Bootdisk` dialog box appears. Insert your second blank floppy disk (mentioned at the beginning of this project). Yes is selected by default, so press **F12** to continue. A message box appears. Press **F12** a second time to begin creating the boot disk. A message box informs you that the disk is being created.
37. After the boot disk is created, the `Lilo Installation` dialog box appears. You see a list of possible locations where the LILO boot manager can be installed. If you have had problems getting Linux to boot after installation, choose the line containing `Master Boot Record`; otherwise, you can choose the line containing `First sector of boot partition`. Use the arrow keys to make your selection, and press **F12** to continue. A second `Lilo Installation` dialog box appears.
38. Some systems require boot parameters in order for Linux to work properly (as described in this chapter). The first time you install Linux you should not assume that these parameters are needed. If they were needed, you could enter them in this dialog box to activate them each time Linux was booted by LILO. Because they are generally not necessary, and should in any case only be added after trying an installation without them, press **F12** to continue.
39. The `Bootable Partitions` dialog box appears. It lists the Red Hat Linux partition and any other partitions containing other operating systems installed on your computer. A Default Boot label column indicates the name assigned to each one. If you need to change the information listed here so that you can use LILO to boot other operating systems, use the arrow keys to select a line containing a partition with another operating system, and then press **Tab** until the Edit button is highlighted. Press **Enter** to open the

Edit Boot Label dialog box. Enter a new label in the Boot label field. When the second LILO configuration dialog box appears again, press **F12** to continue. The Choose a Card dialog box appears.

If the installation program is able to detect information about your video card, you may not see some of the dialog boxes described here.

40. Use the arrow keys to select your video card from the list of hundreds of video cards. On the keyboard, press the first letter in the card's name to move to that point alphabetically in the list. The last item in the list is Unlisted Card, which you can select if the video card you need to configure is not listed. Press **F12** to continue. A message box informs you that the corresponding X server software is being installed on the system. The Monitor Setup dialog box then appears.
41. Select your monitor from the list of monitors shown. If your monitor model is not listed, choose Custom. (In this case, you will be prompted by additional screens to enter your monitor details.) On the keyboard, press a letter key to move to that letter position in the alphabetic listing of monitors. Laptop screens are listed under LCD Panel. When you have selected a monitor, press **F12** to continue. A Screen Configuration message box appears.
42. Review the information in the Screen Configuration message box. For most systems, you can choose Probe to determine your video hardware. When you press **Enter**, a message in a dialog box informs you that the probing is about to begin. Press **F12** to continue. The screen blinks a few times as the installation program searches for information about your video card. Finally, a new dialog box appears informing you that the probing process is finished.
43. Press the **Tab** key repeatedly until the Use Default button is highlighted in white. This causes the installation to use the values determined by probing the video card. Press **F12** to continue. The Starting X message box appears, informing you that the X Window System is about to be tested.
44. Press **F12** to try the X Window System. If the configuration of your video card was successful, you see a graphical message box asking you to click Yes with your mouse to indicate that the graphics mode worked. You can then choose to start your Red Hat Linux system in that graphical mode automatically by clicking Yes in a second graphical message box.
45. The Done dialog box appears, informing you that the installation is finished. Press **F12** to restart the computer in the new Red Hat Linux system.

If you have trouble with the graphical configuration because your video card or monitor is not listed, or the X Window System does not start successfully, choose the Skip option after the X Window System launch fails. You can configure and troubleshoot the graphical system using the information in Chapter 5.

Project 3-3

In this activity you begin to use the new Linux system that you installed in Project 3-2. To complete this activity you should have a newly installed Red Hat Linux system available.

1. Remove any CDs and floppy disks from the computer.
2. Restart the computer. When you see a prompt consisting of the text `boot:`, press **Tab** to see the available operating systems that are installed.
3. The installation program assigned a label (name) to Red Hat Linux during the installation. This label is normally `linux`, which you see when you press **Tab**. Enter this label used for the Red Hat Linux operating system. Review the messages that appear as the system boots.
4. When the boot process is finished, a login screen appears. (This may be a text-mode login screen or a graphical login screen depending on the selections you made in Project 3-2.) Enter the username that you created as a regular user account during the installation.
5. Enter the wrong password. What happens? Can you use the system without entering one of the valid usernames and the corresponding password?
6. Enter the correct username and password of the regular user account to log in to the system.
7. If you are using a text-based screen, enter the command `logout`, or choose **`Logout`** from the main menu of the graphical system if you have entered the graphical desktop automatically.
8. Log in again using the `root` account name and password. Do you see any differences in the appearance of the system when you log in as `root`?

CASE PROJECTS

1. Your friend manages a group of resorts known collectively as the Lakewood Resorts. The Lakewood company has recently started to expand its operations by promoting the resort to vacationers around the country. To support this expansion, they have installed a call center with about 100 computers to handle incoming requests for information and reservations. Each computer will be staffed by a representative who can answer questions and make reservations or send out a resort brochure. All of the computers will be running Linux and will be connected to a large reservations computer located in another office. You have been asked to install Linux on all of the systems. Based on the options you learned in this chapter, would you use a local CD-ROM drive to install each system or place the installation source files on a networked server? Explain your choice. Describe in detail the additional features (beyond those discussed explicitly in the chapter) you would like to see in a Linux distribution or installation program to support your work on this project. Visit the Web sites of several Linux vendors, and see what features you can find that fit your criteria.

2. Assuming you find the features you're looking for to ease the task of installing 100 new systems in a short time, would those features alone lead you to use one Linux distribution over another? What other technical or nontechnical features would be important to you as you prepare to install the call center's computers? Given the standard cost range of Linux distributions mentioned in the chapter, how important is the cost of the Linux product in making your decision?
3. Suppose all the call center's computers already have Microsoft Windows installed. You have been asked to make each one a dual-boot system, so that representatives can use Windows software occasionally if they need to. Does this change the Linux product you would choose? Conduct research on the Internet or through Linux vendors to locate commercial Windows software to help you install (or prepare to install) the Linux systems. Assuming that the Linux systems were already installed and you were later asked to add Windows to each system, how would your arrangement of the Linux partitions (and possibly multiple mount points) affect your ability to make the requested change to the systems? Are some possible future needs too costly to prepare for now?